

Prior-Free Multi-Unit Auctions with Ordered Bidders*

Sayan Bhattacharya[†] Elias Koutsoupias[‡] Janardhan Kulkarni[§] Stefano Leonardi[¶]
Tim Roughgarden^{||} Xiaoming Xu^{**}

Abstract

Prior-free auctions are robust auctions that assume no distribution over bidders' valuations and provide worst-case (input-by-input) approximation guarantees. In contrast to previous work on this topic, we pursue good prior-free auctions with non-identical bidders.

Prior-free auctions can approximate meaningful benchmarks for non-identical bidders only when sufficient qualitative information about the bidder asymmetry is publicly known. We consider digital goods auctions where there is a *total ordering* of the bidders that is known to the seller, where earlier bidders are in some sense thought to have higher valuations. We use the framework of Hartline and Roughgarden (STOC '08) to define an appropriate revenue benchmark: the maximum revenue that can be obtained from a bid vector using prices that are nonincreasing in the bidder ordering and bounded above by the second-highest bid. This *monotone-price benchmark* is always as large as the well-known fixed-price benchmark $\mathcal{F}^{(2)}$, so designing prior-free auctions with good approximation guarantees is only harder. By design, an auction that approximates the monotone-price benchmark satisfies a very strong guarantee: it is, in particular, simultaneously near-optimal for essentially every Bayesian environment in which bidders' valuation distributions have nonincreasing monopoly prices, or in which the distribution of each bidder stochastically dominates that of the next. Even when there is no distribution over bidders' valuations, such an auction still provides a quantifiable input-by-input performance guarantee.

In this paper, we design a simple $O(1)$ -competitive prior-free auction for digital goods with ordered bidders. We also extend the monotone-price benchmark and our $O(1)$ -competitive prior-free auction to multi-unit settings with limited supply.

*Preliminary versions of these results appeared in the Proceedings of the 44th Annual Symposium on Theory of Computing, May 2012, and the Proceedings of the 14th ACM Conference on Electronic Commerce, June 2013.

[†]Max-Planck Institute für Informatik, Saarbrücken. Email: bsayan@mpi-inf.mpg.de.

[‡]University of Oxford and University of Athens. Supported in part by ERC Advanced Grant 321171 (ALGAME). Email: elias@cs.ox.ac.uk.

[§]Duke University. Supported by NSF Awards CCF-1008065 and IIS-0964560. Email: kulkarni@cs.duke.edu.

[¶]Sapienza University of Rome. Supported in part by a Google Faculty Award, EU FET Project MULTIPLEX 317532, and ERC Grant 259515 PAAI. Email: leonardi@dis.uniroma1.it.

^{||}Stanford University. Supported in part by NSF grants CCF-1016885 and CCF-1215965, an ONR PECASE Award, and an AFOSR MURI grant. Email: tim@cs.stanford.edu.

^{**}Duke University. Supported by NSF Awards CCF-1008065 and IIS-0964560. Email: xiaoming@cs.duke.edu.

Report Documentation Page				Form Approved OMB No. 0704-0188	
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE MAY 2014		2. REPORT TYPE		3. DATES COVERED 00-00-2014 to 00-00-2014	
4. TITLE AND SUBTITLE Prior-Free Multi-Unit Auctions with Ordered Bidders				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Stanford University, Department of Computer Science, Stanford, CA, 94305				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT Prior-free auctions are robust auctions that assume no distribution over bidders' valuations and provide worst-case (input-by-input) approximation guarantees. In contrast to previous work on this topic, we pursue good prior-free auctions with non-identical bidders. Prior-free auctions can approximate meaningful benchmarks for non-identical bidders only when sufficient qualitative information about the bidder asymmetry is publicly known. We consider digital goods auctions where there is a total ordering of the bidders that is known to the seller, where earlier bidders are in some sense thought to have higher valuations. We use the framework of Hartline and Roughgarden (STOC '08) to define an appropriate revenue benchmark: the maximum revenue that can be obtained from a bid vector using prices that are nonincreasing in the bidder ordering and bounded above by the second-highest bid. This monotone-price benchmark is always as large as the well-known second-price benchmark $F(2)$, so designing prior-free auctions with good approximation guarantees is only harder. By design an auction that approximates the monotone-price benchmark satisfies a very strong guarantee it is, in particular, simultaneously near-optimal for essentially every Bayesian environment in which bidders' valuation distributions have nonincreasing monopoly prices, or in which the distribution of each bidder stochastically dominates that of the next. Even when there is no distribution over bidders' valuations, such an auction still provides a quantifiable input-by-input performance guarantee. In this paper, we design a simple $O(1)$-competitive prior-free auction for digital goods with ordered bidders. We also extend the monotone-price benchmark and our $O(1)$-competitive prior-free auction to multi-unit settings with limited supply.					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT Same as Report (SAR)	18. NUMBER OF PAGES 18	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

1 Introduction

Suppose you own a set of goods and want to make money by selling them. What is the best way to do it? This question is non-trivial even in *digital goods auctions*, where the seller has an unlimited supply of identical goods (like mp3s), and there are n bidders, each of whom wants only one good and has a private valuation (i.e., maximum willingness-to-pay) for it.

The question becomes easy if the seller has a prior product distribution on bidders' valuations. Since supply is unlimited and valuations are independent, the seller can optimize for each bidder separately. For a bidder i with valuation distribution F_i , the expected revenue is maximized by posting a *monopoly price* — a “take-it-or-leave-it” offer at a price in $\arg\max_p [p \cdot (1 - F_i(p))]$.

What if good prior information is expensive or impossible to acquire? What if a single auction is to be re-used several times, in settings with different or not-yet-known bidder valuations? Are there *prior-free* auctions that admit more robust, “worst-case” revenue guarantees? Particularly germane to this paper, do such auctions exist when none of the bidders are identical?

1.1 Revenue Benchmarks

Goldberg et al. [11, 12] were the first to pursue prior-free auctions, and they proposed a competitive analysis framework based on *revenue benchmarks* [6, 20]; see also the survey by Hartline and Karlin [15].¹ The idea is to define a real-valued function on inputs (i.e., bid vectors) that represents an upper bound on the maximum revenue achievable by any “reasonable” auction on each input. They proposed the *fixed-price benchmark* $\mathcal{F}^{(2)}$ for digital goods auctions, defined as the maximum revenue that can be obtained from a given bid vector by offering every bidder a common posted price that is at most the second-highest bid.

Comparing the revenue of an auction to $\mathcal{F}^{(2)}$ initially looks like an “apples vs. oranges” comparison — the auction does not know bidders' valuations but can employ arbitrary prices, while the benchmark is privy to all the private information but handicapped in the prices it can use. Nevertheless, Goldberg et al. [11] demonstrated the effectiveness of the fixed-price benchmark for meaningful competitive analysis: no auction achieves more than a $\approx .42$ fraction of $\mathcal{F}^{(2)}$ for every bid vector, and there are interesting auctions that obtain a constant fraction of this benchmark on every input.

1.2 The Bayesian Thought Experiment

To extend the revenue benchmark approach to new objective functions and asymmetric outcome spaces, Hartline and Roughgarden [16] advocated a general framework based on a “Bayesian thought experiment.” Roughly, this framework works as follows. The first step is to temporarily think of bidders' valuations as drawn i.i.d. from some valuation distribution. The second step is to characterize the collection \mathcal{C} of all optimal auctions that can arise — those with maximum-possible expected objective function value with respect to some valuation distribution. For example, for revenue maximization in digital goods auctions, \mathcal{C} is the set of common posted prices (bidders are i.i.d. and hence have a common monopoly price). Finally, given a bid vector \mathbf{b} , the benchmark is defined as the maximum objective function value obtained by an auction in \mathcal{C} on the input \mathbf{b} . In digital goods auctions, this is the maximum revenue that can be obtained by offering every bidder a

¹For other recent approaches to the design and analysis of auctions with non-Bayesian sellers, see Chen and Micali [6] and Lopomo et al. [20].

common posted price. Thus, modulo the restriction that prices are at most the second-highest bid, the Bayesian thought experiment automatically regenerates the $\mathcal{F}^{(2)}$ benchmark. (For technical reasons, the upper bound on prices still needs to be added to permit interesting results [11].) More importantly, all benchmarks generated by this framework are automatically well motivated: if the performance of an auction is within a constant factor of such a benchmark for every input, then in particular it is simultaneously near-optimal in every Bayesian i.i.d. environment.² In addition, if there is no distribution over inputs, then the auction still provides a quantifiable input-by-input guarantee.

There are several analogs elsewhere in theoretical computer science: worst-case regret guarantees in online decision-making (e.g., if cost vectors are drawn i.i.d. from a distribution, then the optimal action is time-invariant); and static optimality in data structure design (e.g., if searches are i.i.d., then there is some fixed optimal binary search tree). The framework in [16] and some variants of it have been used successfully to extend the reach of prior-free mechanism design to new objective functions [16] and more complex environments [7, 17, 18].

1.3 Beyond I.I.D. Bidders

The primary goal of this paper is the following.

To design good prior-free auctions for benchmarks derived from non-identical bidders.

Why is this non-trivial? Let’s apply the Bayesian thought experiment to a digital goods auction, now assuming that bidder i ’s valuation is drawn (independently) from its own distribution F_i . For fixed distributions F_1, \dots, F_n , the optimal auction offers each bidder its respective monopoly price. Ranging over all choices of F_1, \dots, F_n , we find that the collection \mathcal{C} corresponds to the set of *all posted price vectors*.³ Thus, for every bid vector \mathbf{b} , there is an auction $A_{\mathbf{b}} \in \mathcal{C}$ that uses the price vector \mathbf{b} and hence obtains the full welfare $\sum_{i=1}^n b_i$ as revenue. There is no digital goods auction that always obtains a constant fraction of the optimal welfare [11], so the Bayesian thought experiment with non-i.i.d. bidders generates a benchmark that is far too strong for meaningful competitive analysis.

The exercise above suggests the following principle for prior-free auction design with non-identical bidders.

Prior-free auctions can approximate benchmarks derived from non-identical bidders only if “sufficient qualitative information” about bidder asymmetry is publicly known.

As an example, suppose there is a publicly known partition of the bidders into groups of otherwise indistinguishable bidders. We then require the Bayesian thought experiment to conform to the public information, meaning that the valuations of bidders in the same group are i.i.d. draws from a distribution. Then, the optimal auctions \mathcal{C} are the price vectors that offer a common posted price to each group of bidders. The induced prior-free benchmark is the maximum revenue that can be obtained from the given bid vector using such a price vector. This is essentially the same benchmark proposed in work on attribute auctions [3, 4] that predates the benchmark framework

²This weaker goal of designing good *prior-independent* auctions — where a distribution over inputs is assumed and used in the analysis of a mechanism, but not in its design — is now studied in its own right. See Dhangwatnotai et al. [9] and the references therein.

³This fact holds even if we restrict the F_i ’s to be, say, uniform distributions with supports $[0, h_i]$ (and hence monopoly prices $h_i/2$).

in [16]. There are prior-free digital goods auctions with expected revenue at least a constant fraction of this benchmark when every group has at least 2 bidders (by an easy reduction to the standard setup) and when there is a constant number of groups [3, 4].

1.4 Ordered Bidders and Stochastic Dominance

What about the general case when all bidders are distinguishable? We initially consider digital goods (unlimited supply) auctions where there is a *total ordering* of the bidders that is known to the seller. Without loss of generality, we assume that bidders are ordered $1, 2, \dots, n$.⁴ Earlier bidders are in some sense expected to have higher valuations. This information could be derived from, for example, zip codes, eBay bidding histories, credit history, previous transactions with the seller, and so on. We emphasize that the known information is only qualitative, and is not quantitative or distributional, as is standard in Bayesian auction design.

To generate a prior-free benchmark, we consider Bayesian thought experiments that conform to the known information. Call the distributions F_1, \dots, F_n *ordered* if the corresponding monopoly prices are nonincreasing. For example, the F_i 's could be:

1. Uniform distributions on intervals $[0, h_i]$ with nonincreasing h_i 's.
2. Exponential distributions with nondecreasing rates.
3. Lognormal distributions with nonincreasing means.

Letting (F_1, \dots, F_n) range over all ordered distributions, the corresponding collection \mathcal{C} of optimal auctions is the set of *monotone* price vectors \mathbf{p} , where $p_1 \geq \dots \geq p_n$. We denote the induced revenue benchmark by $\mathcal{M}^{(1)}$, the maximum revenue that can be obtained from a given bid vector from a monotone price vector. For example, if a bid vector \mathbf{b} is itself monotone, with $b_1 \geq \dots \geq b_n$, then setting $\mathbf{p} = \mathbf{b}$ shows that $\mathcal{M}^{(1)}(\mathbf{b})$ is the full welfare $\sum_{i=1}^n b_i$. If $b_1 \leq \dots \leq b_n$, however, then the revenue-maximizing monotone price vector is simply a constant price — equal to the bid b_i that maximizes $j \cdot b_j$. We emphasize that the benchmark $\mathcal{M}^{(1)}(\mathbf{b})$ is defined, and we demand a good approximation to it, on *every* bid vector \mathbf{b} , including those that defy the semantics of the bidder ordering.

By definition, an auction with revenue at least a constant fraction of $\mathcal{M}^{(1)}$ on every input is simultaneously near-optimal in every Bayesian digital goods auction with independent and ordered distributions. A similar simultaneous approximation result holds under the standard notion of stochastic dominance. Recall that a distribution F_i stochastically dominates another F_{i+1} if $F_i(x) \leq F_{i+1}(x)$ for every $x \geq 0$. Proposition A.1 shows that if F_i stochastically dominates F_{i+1} for every $i = 1, 2, \dots, n-1$, and every distribution is regular⁵, then there is a monotone price vector with expected revenue at least 50% of that of an optimal price vector. It follows that an auction with revenue at least a constant fraction of $\mathcal{M}^{(1)}$ on every input is simultaneously near-optimal in every Bayesian digital goods auction in which the distribution of each bidder stochastically dominates that of the next.

⁴Ties between bidders can also be accommodated easily, either with cosmetic changes to the auction and analysis in this paper, or by handling groups of indistinguishable bidders separately using known techniques.

⁵A distribution F is *regular* [22] if $v - (1 - F(v))/f(v)$ is nondecreasing in v .

1.5 The Monotone Price Benchmark $\mathcal{M}^{(2)}$

Given a digital goods environment with ordered bidders, we define the *monotone price benchmark* $\mathcal{M}^{(2)}(\mathbf{b})$ for every bid vector \mathbf{b} as the maximum revenue obtainable via a monotone price vector in which every price is at most the second-highest bid. As in the standard model with indistinguishable bidders [11], the upper bound on prices is necessary for the existence of prior-free auctions with non-trivial approximation guarantees.⁶ Indeed, since a constant price vector is monotone, $\mathcal{M}^{(2)}(\mathbf{b}) \geq \mathcal{F}^{(2)}(\mathbf{b})$ for every \mathbf{b} and so designing auctions competitive with the monotone-price benchmark is at least as difficult as with the fixed-price benchmark. Taking $b_i = \frac{1}{i}$ for $i = 1, 2, \dots, n$ shows that there exist bid vectors for which $\mathcal{M}^{(2)}(\mathbf{b})$ exceeds $\mathcal{F}^{(2)}(\mathbf{b})$ by an $\Omega(\log n)$ factor. As far as we know, all prior-free auctions proposed prior to the present work are $\Omega(\log n)$ -competitive with $\mathcal{M}^{(2)}$.

The monotone-price benchmark was previously considered, with a completely different motivation, by Aggarwal and Hartline [1]. In [1], which predates the benchmark framework in [16], $\mathcal{M}^{(2)}$ was one of three ad hoc benchmarks proposed for “knapsack auctions,” where bidders have a public size and feasible solutions correspond to subsets of bidders with total size at most a publicly known budget. Aggarwal and Hartline [1] gave a digital goods auction that, for every bid vector \mathbf{b} , has expected revenue at least $\frac{1}{c}\mathcal{M}^{(2)}(\mathbf{b}) - O(h \log \log \log h)$, where $c > 0$ is a constant and h is the ratio between the maximum and minimum bids. Our results improve over those in [1] in several respects: we obtain a constant-factor approximation guarantee without an additive loss term and without any dependence on the magnitude of the valuations, and we also obtain results for limited-supply auctions.

1.6 Our Results: Unlimited Supply

Section 3 gives a digital goods auction that is $O(1)$ -competitive with the monotone price benchmark $\mathcal{M}^{(2)}$. Our auction is simple and natural. It follows the standard approach of randomly partitioning the bidders into two groups, and using one group of bidders to set prices for the other. It computes an optimal monotone price vector for the “training set” of bidders, subject to using prices that are powers of 2, and extends this price vector to the “test set” of bidders. To handle inputs where the monotone price benchmark derives most of its revenue from a small number of bidders, with constant probability we invoke an auction that is $O(1)$ -competitive with the fixed-price benchmark $\mathcal{F}^{(2)}$.

1.7 Our Results: Limited Supply

Section 4 extends our results to multi-unit auctions, where the number of items k can be less than the number of bidders. We consider the analog $\mathcal{M}^{(2,k)}$ of the monotone price benchmark, which maximizes only over (monotone) price vectors that sell at most k units. We prove that every auction that is $O(1)$ -competitive with the benchmark $\mathcal{M}^{(2,k)}$ is simultaneously near-optimal for a range of Bayesian multi-unit environments — roughly, those in which the (ironed) virtual valuation functions of the bidders form a pointwise total ordering. We adapt a reduction from [1] to show

⁶An auction that always has revenue at least a constant fraction of $\mathcal{M}^{(2)}$ is still simultaneously near-optimal in every Bayesian environment with ordered or stochastically dominating distributions, with somewhat worse constant factors, provided these distributions satisfy some mild extra conditions. See Section 4.1 for further discussion.

how to build a limited-supply auction that is $O(1)$ -competitive with respect to $\mathcal{M}^{(2,k)}$ from an unlimited-supply auction that is $O(1)$ -competitive with respect to $\mathcal{M}^{(2)}$.

2 Preliminaries

This section reviews mechanism design basics and digital goods auctions; the expert can skip to Section 3. Section 4 describes the changes needed to accommodate limited-supply settings.

In a *digital goods auction*, there is one seller and n bidders. There is an unlimited supply of identical goods. Each bidder wants only one good, and has a private — i.e., unknown to the seller — *valuation* v_i . We study direct-revelation auctions, in which the bidders report bids \mathbf{b} to the seller, and the seller then decides who wins a good and at what price.⁷ For a fixed (randomized) auction, we use $X_i(\mathbf{b})$ and $P_i(\mathbf{b})$ to denote the winning probability and expected payment of bidder i when the bid profile is \mathbf{b} . As in previous works on prior-free auction design, we consider only auctions that are individually rational — meaning $P_i(\mathbf{b}) \leq b_i \cdot X_i(\mathbf{b})$ for every i and \mathbf{b} — and truthful, meaning that for each bidder i and fixed bids \mathbf{b}_{-i} by the other bidders, bidder i maximizes its quasi-linear utility $v_i \cdot X_i(b_i, \mathbf{b}_{-i}) - P_i(b_i, \mathbf{b}_{-i})$ by setting $b_i = v_i$. Since we consider only truthful auctions, from now on we use bids \mathbf{b} and valuations \mathbf{v} interchangeably.

Truthful and individually rational digital goods auctions have a nice canonical form: for every bidder i there is a (possibly randomized) function $t_i(\mathbf{v}_{-i})$ that, given the valuations \mathbf{v}_{-i} of the other bidders, gives bidder i a “take-it-or-leave-it offer” at the price $t_i(\mathbf{v}_{-i})$. This means that bidder i is given a good if and only if $v_i \geq t_i(\mathbf{v}_{-i})$, in which case it is charged the price $t_i(\mathbf{v}_{-i})$. It is clear that every choice (t_1, \dots, t_n) of such functions defines a truthful, individually rational digital goods auction; conversely, every such auction is equivalent to a choice of (t_1, \dots, t_n) [11]. A special case of such an auction is a *price vector* \mathbf{p} , where each t_i is the constant function $t_i(\mathbf{v}_{-i}) = p_i$. As noted in Section 1, in Bayesian settings with independent valuations, price vectors maximize expected revenue over all truthful and individually rational auctions.

The *revenue* of an auction on the valuation profile \mathbf{v} is the sum of the payments collected from the winners. Let $v^{(2)}$ denote the second-highest valuation of a profile \mathbf{v} . The *fixed-price benchmark* $\mathcal{F}^{(2)}$ is defined, for each valuation profile \mathbf{v} , as the maximum revenue that can be obtained from a constant price vector whose price is at most $v^{(2)}$:

$$\mathcal{F}^{(2)}(\mathbf{v}) = \max_{p \leq v^{(2)}} \left(\sum_{i: v_i \geq p} p \right).$$

Now suppose there is a known ordering on the bidders, say $1, 2, \dots, n$. The *monotone-price benchmark* $\mathcal{M}^{(2)}$ is defined analogously to $\mathcal{F}^{(2)}$, except that non-constant monotone price vectors are also permitted:

$$\mathcal{M}^{(2)}(\mathbf{v}) = \max_{v^{(2)} \geq p_1 \geq p_2 \geq \dots \geq p_n} \left(\sum_{i: v_i \geq p_i} p_i \right). \quad (1)$$

Clearly, $\mathcal{M}^{(2)}(\mathbf{v}) \geq \mathcal{F}^{(2)}(\mathbf{v})$ for every input \mathbf{v} .

We reiterate that the monotonicity and upper-bound constraints are enforced only in the computation of the benchmark $\mathcal{M}^{(2)}$. Auctions, while obviously not privy to the private valuations,

⁷For the questions we ask, the “Revelation Principle” (see, e.g., Nisan [23]) ensures that there is no loss of generality by considering only direct-revelation auctions.

can employ whatever prices they see fit. This is natural for prior-free auctions and also necessary for non-trivial results [10].

Finally, when we say that an auction is α -competitive with or has *approximation factor* α for a benchmark, we mean that the auction’s expected revenue is at least a $1/\alpha$ fraction of the benchmark for every input \mathbf{v} .

3 A Prior-Free $O(1)$ -Approximate Digital Goods Auction with Ordered Bidders

3.1 The Auction

Input: a valuation profile \mathbf{v} for a totally ordered set $N = \{1, 2, \dots, n\}$ of bidders.

1. With probability $1/2$, run a digital goods auction on \mathbf{v} that is $O(1)$ -competitive with respect to the benchmark $\mathcal{F}^{(2)}$, and halt.
2. Choose a subset $A \subseteq N$ uniformly at random, and partition N into the two sets A and $B = N \setminus A$. Let \mathbf{v}^A denote the valuation profile \mathbf{v} in which we set the valuations in B to 0.
3. Compute the revenue-maximizing price vector \mathbf{p} for \mathbf{v}^A that is monotone and that uses prices restricted to be values in $\{2^t : t \in \mathbb{Z}\}$ that are at most the second-highest valuation in \mathbf{v}^A .
4. Sell items to bidders in B only, uses the prices \mathbf{p} .

Figure 1: The auction \mathcal{A}^* .

We propose the auction \mathcal{A}^* , displayed in Figure 1. We next elaborate on the steps of the auction. In the first step, we run an arbitrary digital goods auction that is $O(1)$ -competitive with respect to the fixed-price benchmark $\mathcal{F}^{(2)}$. The best-known approximation factor is 3.12 [19]; there are also very simple auctions with approximation factors 4 [11] and 4.68 [2]. Intuitively, this step is meant to extract good revenue from the set of bidders with valuations almost as high as the second-highest valuation.

The second step of the algorithm randomly partitions the bidders into a “training set” A and a “test set” B . Almost all prior-free auctions have this structure, with the bidders in the training set setting prices for those in the test set. For simplicity, we sell (in the fourth step) only to bidders in the test set B . An obvious optimization is to sell simultaneously to bidders in A , using the bids of B ; this would improve the hidden constant in our approximation guarantee by a factor of 2.

The second step also defines the valuation profile \mathbf{v}^A . This profile has the same length of \mathbf{v} , with the valuations of the bidders in B zeroed out.

The third step computes the monotone price vector that maximizes revenue with respect to the valuation profile \mathbf{v}^A , subject to the extra constraint that every price is a (possibly negative) integer power of 2 bounded above by the second-highest valuation of \mathbf{v}^A . Using dynamic programming, this step (and hence the entire auction) can be implemented in polynomial time.

Let \mathbf{p} be the price vector computed in the third step. In the language of Section 2, the fourth step sets the take-it-or-leave-it offer $t_i(\mathbf{v}_{-i})$ to $+\infty$ for bidders $i \in A$ and to p_i for bidders $i \in B$.

Since \mathbf{p} is computed using only the valuations of the bidders in A , these $t_i(\mathbf{v}_{-i})$'s are well defined and the auction \mathcal{A}^* is truthful and individually rational.

We prove the following.

Theorem 3.1. *There is a constant $c > 0$ such that, for every input \mathbf{v} , the expected revenue of the auction \mathcal{A}^* is at least $c \cdot \mathcal{M}^{(2)}(\mathbf{v})$.*

Very roughly, the intuition behind the auction \mathcal{A}^* and Theorem 3.1 is the following. Consider first a valuation profile \mathbf{v} in which a constant fraction of the revenue in $\mathcal{M}^{(2)}(\mathbf{v})$ is provided by (a constant number of) bidders with valuation at least a constant times $\mathcal{M}^{(2)}(\mathbf{v})$. In this case, the fixed-price benchmark $\mathcal{F}^{(2)}(\mathbf{v})$ is within a constant factor of $\mathcal{M}^{(2)}(\mathbf{v})$, and the first step of \mathcal{A}^* ensures that the auction is constant-competitive. Thus, the difficult inputs are those in which a large number of bidders contribute to $\mathcal{M}^{(2)}(\mathbf{v})$. For these inputs, however, concentration bounds should imply that the test set A strongly resembles the training set B , and hence the computed price vector \mathbf{p} should generalize well. We note, however, that this high-level intuition appears also in previous works [1] that obtained worse bounds; to prove a constant-competitive guarantee, the analysis has to be executed with some care.

3.2 Analysis Preliminaries: Some Important Events

This section identifies some important probabilistic events and proves that, for every valuation profile, they hold with constant probability over the coin flips of \mathcal{A}^* (i.e., the random choice of A). The next section shows that the revenue of \mathcal{A}^* is close to $\mathcal{M}^{(2)}(\mathbf{v})$ whenever these events hold, which implies Theorem 3.1.

For the rest of this section, fix an arbitrary valuation profile \mathbf{v} . Let $\text{REV}^A(\mathbf{p})$ denote the revenue extracted by the prices \mathbf{p} from the bidders in A in the third step of \mathcal{A}^* . Let $\text{REV}^B(\mathbf{p})$ denote the revenue extracted by \mathbf{p} from the bidders in B in the fourth step of \mathcal{A}^* . Define $\text{REV}(\mathbf{p}) = \text{REV}^A(\mathbf{p}) + \text{REV}^B(\mathbf{p})$.

By definition, event \mathcal{E}_1 occurs when $\text{REV}(\mathbf{p}) \geq \mathcal{M}^{(2)}(\mathbf{v})/6$.

Lemma 3.2. *The event \mathcal{E}_1 holds with probability at least $1/16$.*

Proof. Let \mathbf{p}^* achieve the maximum in (1) for \mathbf{v} . With probability $1/4$, the bidders with the highest and second-highest valuations lie in A . Given this event, the conditional expected revenue from bidders in A under the price vector \mathbf{p}^* is at least $\mathcal{M}^{(2)}(\mathbf{v})/2$. By Markov's inequality, the conditional expected revenue from bidders in A under \mathbf{p}^* is at least $\mathcal{M}^{(2)}(\mathbf{v})/3$ with probability at least $\frac{1}{4}$. Since the bidders with highest and second-highest valuations lie in A , rounding every price of \mathbf{p}^* down to the nearest power of 2 yields a candidate for the price vector \mathbf{p} computed by the auction \mathcal{A}^* in its third step, and the revenue extracted by this candidate is at least half that of \mathbf{p}^* . Thus, $\text{REV}^A(\mathbf{p}) \geq \mathcal{M}^{(2)}(\mathbf{v})/6$ with probability at least $\frac{1}{4} \cdot \frac{1}{4} = \frac{1}{16}$. Since $\text{REV}(\mathbf{p}) \geq \text{REV}^A(\mathbf{p})$ with probability 1, the lemma follows. \square

Identifying the next collection of important events requires some definitions.

Definition 3.1. *For every integer $l \geq 0$, the l -th price level is the (unique) price q in $\{2^t : t \in \mathbb{Z}\}$ such that $\mathcal{M}^{(2)}(\mathbf{v})/2^{l+1} < q \leq \mathcal{M}^{(2)}(\mathbf{v})/2^l$.*

We use $p_{(l)}$ to denote the l -th price level.

Definition 3.2. For a nonnegative integer l , a level- l triple (i, j, l) is two bidders $i < j$ with $v_i, v_j \geq p_{(l)}$.

We denote by W_{ijl} the bidders between i and j (inclusive) that would win at a price of $p_{(l)}$:

$$W_{ijl} = \{k \in N : i \leq k \leq j \text{ and } v_k \geq p_{(l)}\}.$$

We call a level- l triple (i, j, l) *large* if $|W_{ijl}| \geq 288l$. We call a level- l triple (i, j, l) *balanced* under a partition of the bidders into A and B if its winners are split $\frac{1}{3}$ - $\frac{2}{3}$ or better between the two sets:

$$\frac{1}{3} \times |W_{ijl}| \leq |A \cap W_{ijl}|, |B \cap W_{ijl}| \leq \frac{2}{3} \times |W_{ijl}|.$$

By definition, event $\mathcal{E}_2(l)$ occurs when every large level- l -triple is balanced. We let \mathcal{E}_2 denote $\cap_{l \geq 24} \mathcal{E}_2(l)$. We proceed to lower bound the probability of this event.

Claim 3.1. For every integer $l \geq 0$, the number of level- l -triples is at most 2^{2l+2} .

Proof. Consider a bidder k with valuation $v_k \geq p_{(l)} > \mathcal{M}^{(2)}(\mathbf{v})/2^{l+1}$. The definition of $\mathcal{M}^{(2)}(\mathbf{v})$ implies that there are at most 2^{l+1} such bidders. Since a level- l -triple (i, j, l) is uniquely determined by two bidders with valuation at least $p_{(l)}$, there are at most $(2^{l+1})^2 = 2^{2l+2}$ level- l -triples. \square

We use the following version of the Chernoff bound (see e.g. [21]).

Theorem 3.3. Let T_1, \dots, T_m be i.i.d random variables such that $T_i \in \{0, 1\}$ for all $i \in \{1, \dots, m\}$. Define $T = \sum_{i=1}^m T_i$ and $\mu = E[T]$. For all $0 < \delta < 1$:

$$\Pr[(1 - \delta)\mu \leq T \leq (1 + \delta)\mu] \geq 1 - 2 \times \exp\left(-\frac{\mu\delta^2}{4}\right).$$

Claim 3.2. For every $l \geq 24$, $\Pr[\mathcal{E}_2(l)] \geq 1 - 1/2^l$.

Proof. Fix a large level- l -triple (i, j, l) . By definition, the number of winning bidders in (i, j, l) is at least $288l$. Since each of these bidders is included in the set A independently and uniformly at random, Theorem 3.3 implies that the triple (i, j, l) is not balanced with probability at most $2/e^{4l}$. By Claim 3.1, there are at most 2^{2l+2} level- l -triples. By the union bound, the probability that some large level- l -triple is not balanced is at most $2^{2l+2} \times 2/e^{4l} \leq 1/2^l$ when $l \geq 24$. \square

Lemma 3.4. The event \mathcal{E}_2 holds with probability at least $31/32$.

Proof. By Claim 3.2 and the union bound,

$$1 - \Pr[\mathcal{E}_2] \leq \sum_{l \geq 24} (1 - \Pr[\mathcal{E}_2(l)]) \leq \sum_{l \geq 24} \frac{1}{2^l} \leq \frac{1}{32}.$$

\square

Lemmas 3.2 and 3.4 imply the following.

Corollary 3.5. $\Pr[\mathcal{E}_1 \cap \mathcal{E}_2] \geq 1/32$.

3.3 The Main Analysis

Fix a valuation profile \mathbf{v} . Let $I_l(\mathbf{p}) = \{j \in N : p_j = p_{(l)}\}$ denote the bidders offered the price $p_{(l)}$ in \mathbf{p} . Since \mathbf{p} is a monotone price vector, $I_l(\mathbf{p})$ is an interval of bidders. Let $W_l(\mathbf{p}) = \{i \in I_l(\mathbf{p}) : v_i \geq p_{(l)}\}$ denote the bidders of $I_l(\mathbf{p})$ that win under the price vector \mathbf{p} . The interval $I_l(\mathbf{p})$ is *good* if $|W_l(\mathbf{p})| \geq 288l$ and *bad* otherwise. Let $\text{REV}_l(\mathbf{p}) = |W_l(\mathbf{p})| \times p_{(l)}$ denote the contribution of these bidders toward $\text{REV}(\mathbf{p})$. Since every bidder belongs to exactly one interval, $\text{REV}(\mathbf{p}) = \sum_{l \geq 0} \text{REV}_l(\mathbf{p})$.

The next claim shows that the bad intervals $I_l(\mathbf{p})$ with $l \geq 24$ contribute relatively little revenue.

Claim 3.3. *With probability 1,*

$$\sum_{l \geq 24 : I_l(\mathbf{p}) \text{ is bad}} \text{REV}_l(\mathbf{p}) \leq \frac{1}{18} \times \mathcal{M}^{(2)}(\mathbf{v}).$$

Proof. Fix a bad interval $I_l(\mathbf{p})$. Since $|W_l(\mathbf{p})| < 288l$ and $p_{(l)} \leq \mathcal{M}^{(2)}(\mathbf{v})/2^l$,

$$\text{REV}_l(\mathbf{p}) = |W_l(\mathbf{p})| \times p_{(l)} < \frac{288l}{2^l} \times \mathcal{M}^{(2)}(\mathbf{v}).$$

Summing over all bad intervals $I_l(\mathbf{p})$ with $l \geq 24$, we obtain

$$\sum_{l \geq 24 : I_l(\mathbf{p}) \text{ is bad}} \text{REV}_l(\mathbf{p}) \leq \sum_{l \geq 24} \frac{288l}{2^l} \times \mathcal{M}^{(2)}(\mathbf{v}) \leq \frac{1}{18} \times \mathcal{M}^{(2)}(\mathbf{v}).$$

□

We can now prove our main result.

Proof of Theorem 3.1: Fix a valuation profile \mathbf{v} . First suppose that $\mathcal{M}^{(2)}(\mathbf{v}) \leq 432 \cdot \mathcal{F}^{(2)}(\mathbf{v})$. With 50% probability, the auction \mathcal{A}^* executes an auction that is α -competitive with $\mathcal{F}^{(2)}$ for a constant α . Thus, the expected revenue of \mathcal{A}^* on this input is at least $\mathcal{F}^{(2)}(\mathbf{v})/2\alpha \geq \mathcal{M}^{(2)}(\mathbf{v})/864\alpha$.

For the rest of the proof, we assume that $\mathcal{M}^{(2)}(\mathbf{v}) > 432 \cdot \mathcal{F}^{(2)}(\mathbf{v})$. We claim that in this case, with probability 1, the first few intervals contribute little revenue:

$$\sum_{l=0}^{23} \text{REV}_l(\mathbf{p}) \leq \mathcal{M}^{(2)}(\mathbf{v})/18. \quad (2)$$

For otherwise, there is an interval $I_h(\mathbf{p})$ with $h \in [0, 23]$ with

$$\text{REV}_h(\mathbf{p}) = |W_h(\mathbf{p})| \times p_{(h)} > \mathcal{M}^{(2)}(\mathbf{v})/(18 \times 24).$$

Consider the fixed-price vector \mathbf{p}' with common offer price $p_{(h)}$. Since every price of \mathbf{p} is at most the second-highest valuation in \mathbf{v}^A (and hence in \mathbf{v}), the same holds for \mathbf{p}' . The fixed-price benchmark $\mathcal{F}^{(2)}(\mathbf{v})$ is at least the revenue extracted by \mathbf{p}' , which is at least $|W_h(\mathbf{p})| \times p_{(h)} > \mathcal{M}^{(2)}(\mathbf{v})/432$; this contradicts our initial assumption.

Assume now that $\mathcal{E}_1 \cap \mathcal{E}_2$ holds. Since \mathcal{E}_1 holds, $\text{REV}(\mathbf{p}) \geq \mathcal{M}^{(2)}(\mathbf{v})/6$. Combining this with Claim 3.3 and (2), the good intervals from the bigger levels provide large revenue:

$$\sum_{l \geq 24 : I_l(\mathbf{p}) \text{ is good}} \text{REV}_l(\mathbf{p}) \geq \left(\frac{1}{18}\right) \times \mathcal{M}^{(2)}(\mathbf{v}). \quad (3)$$

Consider a good interval $I_l(\mathbf{p})$ with $l \geq 24$. Denote by i and j the first and last bidders in $W_l(\mathbf{p})$, respectively, so that for all $k \in W_l(\mathbf{p})$ we have $i \leq k \leq j$. Since $p_i = p_j = p_{(l)}$ and $v_i, v_j \geq p_{(l)}$, (i, j, l) is a level- l -triple. Because the interval is good, $|W_l(\mathbf{p})| \geq 288l$, and hence the triple (i, j, l) is large. Since \mathcal{E}_2 holds and $l \geq 24$, the triple (i, j, l) is balanced. Hence,

$$|W_l(\mathbf{p}) \cap B| \geq \left(\frac{1}{3}\right) \times |W_l(\mathbf{p})|$$

and the revenue from the bidders in $I_l(\mathbf{p}) \cap B$ under \mathbf{p} is at least $(1/3) \times \text{REV}_l(\mathbf{p})$. Summing over all good intervals $I_l(\mathbf{p})$ with $l \geq 24$ and applying (3) yields

$$\text{REV}^B(\mathbf{p}) \geq \sum_{l \geq 24: I_l(\mathbf{p}) \text{ is good}} \left(\frac{1}{3}\right) \times \text{REV}_l(\mathbf{p}) \geq \left(\frac{1}{54}\right) \times \mathcal{M}^{(2)}(\mathbf{v}). \quad (4)$$

Since the auction \mathcal{A}^* executes steps 2–4 with 50% probability, and since $\Pr[\mathcal{E}_1 \cap \mathcal{E}_2] \geq 1/32$ (Corollary 3.5), the expected revenue of \mathcal{A}^* on such an input \mathbf{v} is at least

$$\frac{1}{2} \times \frac{1}{32} \times \mathbf{E}[\text{REV}^B(\mathbf{p}) \mid \mathcal{E}_1 \cap \mathcal{E}_2] \geq \frac{\mathcal{M}^{(2)}(\mathbf{v})}{3456}.$$

This completes the proof. ■

4 Limited-Supply Multi-Unit Auctions

This section extends our results to multi-unit auctions with limited supply. To develop this theory, we extend the monotone price benchmark $\mathcal{M}^{(2)}$ to the case of an arbitrary number $k \geq 2$ of units for sale. We call a price vector \mathbf{p} *feasible* for the valuation profile \mathbf{v} and supply limit k if: (i) $p_1 \geq p_2 \geq \dots \geq p_n$; (ii) all prices are at most the second-highest valuation of \mathbf{v} ; and (iii) there are at most k bidders i with $v_i > p_i$. We allow our benchmark to break ties in an optimal way. Thus, the revenue earned by a feasible price vector is $\sum_{i: v_i > p_i} p_i$ plus, if there are ℓ items remaining after these sales, the sum of the prices offered to up to ℓ bidders i with $v_i = p_i$. We define the k -unit monotone price benchmark $\mathcal{M}^{(2,k)}(\mathbf{v})$ as the maximum revenue obtained by a price vector that is feasible for \mathbf{v} and k .

There are two main issues to address. The first issue is to identify a class of prior distributions such that approximating $\mathcal{M}^{(2,k)}$ pointwise implies simultaneous approximation of the optimal expected revenue across all Bayesian multi-unit settings with priors belonging to the class. The challenge, relative to the unlimited-supply setting in Section 3, is that limited-supply Bayesian optimal auctions are considerably more complex than unlimited-supply ones. Section 4.1 shows, essentially, that the benchmark $\mathcal{M}^{(2,k)}(\mathbf{v})$ is meaningful whenever the valuation distributions have pointwise ordered ironed virtual valuations. The second issue is to design auctions competitive with the benchmark $\mathcal{M}^{(2,k)}(\mathbf{v})$. Section 4.2 accomplishes this by adapting a reduction in [1] to show how to obtain a limited-supply auction that is $O(1)$ -competitive with respect to $\mathcal{M}^{(2,k)}(\mathbf{v})$ from a digital goods auction that is $O(1)$ -competitive with respect to $\mathcal{M}^{(2)}$.

4.1 Justifying the k -Unit Monotone Price Benchmark

The goal of this section is to prove that every prior-free auction that is $O(1)$ -competitive with the benchmark $\mathcal{M}^{(2,k)}(\mathbf{v})$ has expected revenue at least a constant fraction of optimal in every

Bayesian multi-unit environment with valuation distributions lying in a prescribed class. Making this precise requires some terminology and facts from the theory of Bayesian optimal auction design, as developed by Myerson [22].

4.1.1 Optimal Auction Theory

Consider a bidder with valuation drawn from a prior distribution F with positive and continuous density f on some interval. The *virtual value* v at a point v in the support is defined as

$$\phi(v) = v - \frac{1 - F(v)}{f(v)}.$$

For example, if F is the uniform distribution on $[0, b]$, then the corresponding virtual valuation function is $\phi(v) = 2v - b$.

For clarity, we first discuss the case of *regular* distributions, meaning distributions with nondecreasing virtual valuation functions. In this case, the Bayesian optimal auction awards items to the (at most k) bidders with the highest positive virtual valuations. The payment of a winning bidder is the minimum bid at which it would continue to win (keeping others' bids the same). That is, if the $(k + 1)$ th highest virtual valuation is z , then every winning bidder i pays $\phi_i^{-1}(\max\{0, z\})$. For these prices to be related to the monotone price benchmark, we need to impose conditions on the $\phi_i^{-1}(z)$'s. This contrasts with the unlimited-supply setting, where restricting the $\phi_i^{-1}(0)$'s — that is, the monopoly reserve prices — to be nonincreasing in i is enough to justify the monotone-price benchmark (Section 1.4). Since the $(k + 1)$ th highest virtual valuation could be anything, the natural requirement is to restrict $\phi_i^{-1}(z)$ to be nonincreasing in i for every non-negative number z .

Accommodating irregular distributions, for which the optimal Bayesian auction is more complicated, presents additional complications. Each virtual valuation function ϕ_i is replaced by the “nearest nondecreasing approximation”, called the *ironed virtual valuation function* $\bar{\phi}_i$. The optimal auction awards the items to the (at most k) bidders with the highest positive ironed virtual valuations. Since ironed virtual valuation functions typically have non-trivial constant regions, ties can occur, and we assume that ties are broken randomly. That is, if there are k items, a group S of bidders that all have ironed virtual valuation equal to $z > 0$, and $\ell < k$ bidders with ironed virtual value greater than z with $\ell + |S| > k$, then $k - \ell$ winners from S are chosen uniformly at random.

4.1.2 Pointwise Ordered Distributions

We call valuation distributions F_1, \dots, F_n *pointwise ordered* if $\bar{\phi}_i^{-1}(z)$ is nonincreasing in i for every non-negative z .⁸ The motivating parametric examples discussed in Section 1.4 — uniform distributions with intervals $[0, h_i]$ and nonincreasing h_i 's, exponential distributions with nondecreasing rates, and lognormal distributions with nonincreasing means — are pointwise ordered in this sense.

We also require a second condition, which we inherit from the standard i.i.d. unlimited-supply setting. The issue is that, with arbitrary irregular distributions, no prior-free auction can be simultaneously near-optimal in all Bayesian environments, even with i.i.d. bidders and unlimited supply.⁹

⁸Since $\bar{\phi}_i$ is continuous and nondecreasing, $\bar{\phi}_i^{-1}(z)$ is an interval. If the inverse image has multiple points, we define $\bar{\phi}_i^{-1}(z)$ by the infimum. If the inverse image is empty, we define $\bar{\phi}_i^{-1}(z)$ as the left or right endpoint of the distribution's support, as appropriate.

⁹Informally, consider valuation distributions that take on only two values, one very large (say M) and the other 0. Suppose the probability of having a very large valuation is very small (say $1/n^2$). If the distribution is known, the

Various mild conditions are sufficient to rule out this problem; see [16] for a discussion. Here, for simplicity, we restrict attention to *well-behaved* Bayesian multi-unit environments, meaning that the Bayesian optimal auction derives at most a constant fraction (90%, say) of its revenue from outcomes in which some winner is charged a price higher than the second-highest valuation. (Such a winner is necessarily the bidder with the highest valuation.) Textbook distributions generally yield well-behaved environments.

4.1.3 Connecting $\mathcal{M}^{(2,k)}$ to Bayesian Multi-Unit Settings

The main result of this section is that approximating the k -unit monotone price benchmark $\mathcal{M}^{(2,k)}$ guarantees simultaneous approximation of the optimal auction in all well-behaved Bayesian multi-unit environments with pointwise ordered distributions. We require the following lemma, which states that “projecting” onto a subset of bidders can only decrease the value of the benchmark $\mathcal{M}^{(2,k)}$.

Lemma 4.1. *For every valuation profile \mathbf{v} , $k \geq 2$, and subset S of the bidders with induced profile \mathbf{v}^S , $\mathcal{M}^{(2,k)}(\mathbf{v}) \geq \mathcal{M}^{(2,k)}(\mathbf{v}^S)$.*

Proof. Fix an input \mathbf{v} , with monotone prices \mathbf{p}^* determining $\mathcal{M}^{(2,k)}(\mathbf{v})$. By induction, we only need to show that adding a single new bidder i to an arbitrary position in the ordering can only increase the value of the benchmark. Start by offering i the same price r as its predecessor in the ordering (or the second-highest valuation, if there is no predecessor). If i rejects (i.e., $v_i < r$), this extended price vector is feasible and we are done (the optimal feasible price vector is only better). If i accepts (i.e., $v_i \geq r$) and the price vector is infeasible (with $k + 1$ winners), then we argue as follows. Go through the bidders after i one by one, increasing the offer price to r . This preserves monotonicity. If a previously winning bidder ever rejects this higher offer price, we are done — feasibility is restored and the overall revenue is higher. If not, there is now a “suffix” of bidders with the common offer price r . We now increase the common offer price to the bidders in this suffix until it equals the price offered to the previous bidder in \mathbf{p}^* . This increases the number of bidders in the suffix, and the price-increasing process continues. Eventually a bidder that was winning under \mathbf{p}^* will reject the new offer price — otherwise we contradict the optimality of \mathbf{p}^* . This leaves us with a feasible monotone price vector with revenue at least that of the original one. \square

Theorem 4.2. *If the expected revenue of a multi-unit auction \mathcal{A} is at least a constant fraction of $\mathcal{M}^{(2,k)}(\mathbf{v})$ on every input, then, in every well-behaved multi-unit Bayesian environment with pointwise ordered distributions, the expected revenue of \mathcal{A} is at least a constant fraction of that of the optimal auction for the environment.*

Proof. Fix an auction \mathcal{A} that is β -competitive with $\mathcal{M}^{(2,k)}$. Fix a well-behaved Bayesian multi-unit environment with pointwise ordered valuation distributions F_1, \dots, F_n . Let OPT be the optimal auction for this environment. We claim that, for every input \mathbf{v} in which the revenue collected by OPT from the bidder with the highest valuation is at most the second-highest valuation, the benchmark $\mathcal{M}^{(2,k)}(\mathbf{v})$ is at least half the expected revenue of OPT on \mathbf{v} . This implies that the expected revenue of \mathcal{A} is at least $1/2\beta$ times that of OPT on this input. Since the environment is well behaved, the theorem follows from this claim.

optimal auction uses a reserve price of M for each bidder. Elementary arguments, as in [16], show that no single auction is near-optimal for all values of M .

To prove the claim, fix an input \mathbf{v} , as above. Recall that OPT , as a Bayesian optimal auction, awards items to the (at most k) bidders with the highest positive ironed virtual valuations, breaking ties randomly. The tricky case of the proof is when ties occur. Assume there are k items, a group S of bidders with common ironed virtual value $z > 0$, and a group T of $\ell \in (k - |S|, k)$ bidders with ironed virtual value greater than z (so $|S| > k - \ell$). We next explicitly compute the payments collected by OPT on this input, using the standard payment formula for incentive-compatible mechanisms (see [22] or [14]). Let a_i and b_i denote the left and right endpoints, respectively, of the interval of values v that satisfy $\bar{\phi}_i(v) = z$. Since the distributions are pointwise ordered, the a_i 's and the b_i 's are nonincreasing in i . Let $q = (k - \ell)/|S|$ denote the winning probability of a bidder in S . Define $q' = (k - \ell + 1)/(|S| + 1)$ as the hypothetical winning probability of a bidder in T if it lowered its bid to the value $\bar{\phi}_i^{-1}(z)$. The expected payment of a bidder i in S is $qa_i - a_i$ in the event that it wins. The payment of a bidder i in T (who wins with certainty) is $q'a_i + (1 - q')b_i$. To complete the proof, we argue that $\mathcal{M}^{(2,k)}(\mathbf{v})$ is at least the revenue collected by OPT from the bidders in S , and also at least that from the bidders in T .

Recall from Lemma 4.1 that projecting onto a subset of bidders only decreases the value of $\mathcal{M}^{(2,k)}(\mathbf{v})$. First, project onto the k bidders of S with the highest a_i values. Consider charging each such bidder the price a_i . This is a monotone price vector. By our assumption on the input \mathbf{v} , all of these prices are at most the second-highest valuation in \mathbf{v} . By the definitions, $v_i \geq a_i$ for every bidder $i \in S$ so every offer will be accepted. The resulting revenue is at least the revenue collected by OPT from bidders in S , and $\mathcal{M}^{(2,k)}(\mathbf{v})$ is only higher.

Similarly, project onto the (at most k) bidders of T , and consider charging each such bidder i the price $q'a_i + (1 - q')b_i$. Again, this is a monotone price vector with all prices bounded above by the second-highest valuation of \mathbf{v} , and every offer will be accepted. The value of the monotone price benchmark can only be larger, so $\mathcal{M}^{(2,k)}(\mathbf{v})$ is also at least the revenue collected by OPT from bidders in T . The proof is complete. \square

4.2 Reduction from Limited to Unlimited Supply

Having justified the k -unit monotone price benchmark $\mathcal{M}^{(2,k)}(\mathbf{v})$, we turn to designing auctions that approximate it well. We show that competing with this benchmark reduces to competing with the benchmark $\mathcal{M}^{(2)}$ in unlimited-supply settings. The reduction from limited to unlimited supply for ordered bidders was given in [1] for knapsack auctions. This reduction is also a generalization of the one in [11] for identical bidders. The idea is to first identify the k “most valuable” bidders, and then run an unlimited-supply auction on them. Observe that the most valuable bidders with an ordering are not necessarily those with the highest valuations. For example, a high-valuation bidder late in the ordering need not be valuable, because extracting high revenue from it might necessitate excluding many moderate-valuation bidders earlier in the ordering. We report the “black-box reduction” of [1] in Figure 2.

Theorem 4.3. *If \mathcal{A} is a truthful unlimited-supply auction with ordered bidders that is β -competitive with respect to $\mathcal{M}^{(2)}$, then the BLACK-BOX REDUCTION (BBR) auction is a truthful limited-supply auction with ordered bidders that is 2β -competitive with $\mathcal{M}^{(2,k)}(\mathbf{v})$.*

Proof. The analysis in [1] immediately implies that the BLACK-BOX REDUCTION (BBR) auction is truthful, individually rational, and has at most k winners. We also note that the first step can be implemented efficiently using dynamic programming, so if \mathcal{A} runs in polynomial time, then so does the BLACK-BOX REDUCTION (BBR) auction.

Input: A valuation profile \mathbf{v} for a totally ordered set $N = \{1, 2, \dots, n\}$ of bidders and k identical items. \mathcal{A} denotes a truthful digital goods (unlimited-supply) auction with ordered bidders.

1. Let \mathbf{p} achieve the optimum monotone price benchmark $\mathcal{M}^{(2,k)}(\mathbf{v})$ for \mathbf{v} and k . Let $S = \{i \in N : v_i \geq p_i\}$ be the set of winners under \mathbf{p} .
2. Run the unlimited supply auction \mathcal{A} on the bidders S , with the induced bidder ordering.
3. Charge suitable prices so that truthful reporting is a dominant strategy for every bidder.

Figure 2: The auction BLACK-BOX REDUCTION (BBR).

We prove the performance guarantee by arguing two statements: (i) the unlimited supply benchmark $\mathcal{M}^{(2)}(\mathbf{v}^S)$ applied to S is at least half of the limited-supply benchmark $\mathcal{M}^{(2,k)}(\mathbf{v})$ applied to the original bidder set; and (ii) the expected revenue of BLACK-BOX REDUCTION (BBR) on the original bidder set is at least that of the auction \mathcal{A} with the bidders S . The second statement follows immediately from the facts that the winners of BLACK-BOX REDUCTION (BBR) are the same as those of \mathcal{A} , and that the winners' payments are only higher. For statement (i), consider prices \mathbf{p} that determine the benchmark $\mathcal{M}^{(2,k)}(\mathbf{v})$. The projection \mathbf{p}^S of this price vector onto the set S of bidders has revenue exactly $\mathcal{M}^{(2,k)}(\mathbf{v})$. If \mathbf{p}^S is feasible, then it certifies that the benchmark $\mathcal{M}^{(2)}(\mathbf{v}^S)$ is at least $\mathcal{M}^{(2,k)}(\mathbf{v})$. The only issue is if \mathbf{p}^S uses a price larger than the second-highest valuation $v^{(2,S)}$ of v^S . Setting $\hat{p}_i = \min\{p_i, v^{(2,S)}\}$ for each $i \in S$ yields a monotone and feasible price vector $\hat{\mathbf{p}}$. Every price of \mathbf{p}^S is at most the second-highest valuation $v^{(2)}$ of the original bidders, and \mathbf{p}^S extracts a price higher than $v^{(2,S)}$ from at most one bidder of S (the one with highest valuation). Thus, the revenue extracted by $\hat{\mathbf{p}}^S$ from \mathbf{v}^S is at least that of \mathbf{p}^S , less $v^{(2)}$. Since $\mathcal{M}^{(2,k)}(\mathbf{v}) \geq 2v^{(2)}$ — consider the price vector that offers $v^{(2)}$ to everybody — $\hat{\mathbf{p}}^S$ retains at least half the revenue of \mathbf{p}^S . Statement (i) and the theorem follow. \square

Of course, we can use the auction \mathcal{A}^* from Section 3 in Theorem 4.3 to obtain a truthful limited-supply auction that is $O(1)$ -competitive with the benchmark $\mathcal{M}^{(2,k)}(\mathbf{v})$. Theorem 4.2 implies that the resulting auction enjoys a strong simultaneous approximation guarantee in Bayesian environments.

5 Conclusions

This paper introduced the problem of prior-free auction design with ordered bidders. The bidder ordering represents qualitative information about which bidders are in some sense expected to have higher valuations. We used the “Bayesian thought experiment” of [16] to prove that every auction that is $O(1)$ -competitive with the monotone-price benchmark $\mathcal{M}^{(2)}$ of [1] is simultaneously near-optimal across a wide range of Bayesian settings. Our main result is a construction of such a prior-free auction. We also extend the monotone price benchmark, its connection to Bayesian auction design, and our $O(1)$ -competitive prior-free auction to limited-supply settings.

There are a number of promising directions for future research.

1. For the problems studied in this paper, it would be interesting to design auctions with much

better constant-factor approximation guarantees. The profit-extraction and consensus techniques, as in [13], could be useful for this purpose.

2. For settings more general than identical goods, it would be interesting to generalize all of the contributions of this paper — the prior-free benchmark, the connection to Bayesian settings, and the design of $O(1)$ -competitive auctions. Matroid settings [18], where the feasible outcomes correspond to independent sets of a matroid on the bidder set, are a natural place to begin.
3. It would be interesting to incorporate budgets into the model. Thus far, all work on prior-free auction design with budgets handles only equal budgets [8]. Can any of our techniques for heterogeneous (ordered) bidders be transferred to deal with heterogeneous budgets?
4. Finally, it would be interesting to pursue prior-independent guarantees in the spirit of [9] in Bayesian environments with ordered or stochastically dominating distributions.

References

- [1] G. Aggarwal and J. D. Hartline. Knapsack auctions. In *Proceedings of the 17th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1083–1092, 2006.
- [2] S. Alaei, A. Malekian, and A. Srinivasan. On random sampling auctions for digital goods. In *Proceedings of the 10th ACM Conference on Electronic Commerce (EC)*, pages 187–196, 2009.
- [3] M.-F. Balcan, A. Blum, J. Hartline, and Y. Mansour. Mechanism Design via Machine Learning. *Journal of Computer and System Sciences*, 74(8):1245–1270, 2008.
- [4] A. Blum and J. D. Hartline. Near-optimal online auctions. In *Proceedings of the 16th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1156–1163, 2005.
- [5] J. Bulow and P. Klemperer. Auctions versus negotiations. *American Economic Review*, 86(1):180–194, 1996.
- [6] J. Chen and S. Micali. Mechanism design with set-theoretic beliefs. In *Proceedings of the 52nd Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 87–96, 2011.
- [7] N. Devanur and J. D. Hartline. Limited and online supply and the Bayesian foundations of prior-free mechanism design. In *Proceedings of the 10th ACM Conference on Electronic Commerce (EC)*, pages 41–50, 2009.
- [8] Nikhil R. Devanur, Bach Q. Ha, and Jason D. Hartline. Prior-free auctions for budgeted agents. In *Proceedings of the 14th ACM Conference on Electronic Commerce (EC)*, pages 287–304, 2013.
- [9] Peerapong Dhangwatnotai, Tim Roughgarden, and Qiqi Yan. Revenue maximization with a single sample. In *Proceedings of the 11th ACM Conference on Electronic Commerce (EC)*, pages 129–138, 2010.
- [10] A. Goldberg and J. Hartline. Envy-free auctions for digital goods. In *Proceedings of the 4th ACM Conference on Electronic Commerce (EC)*, pages 29–35, 2003.

- [11] A. V. Goldberg, J. D. Hartline, A. Karlin, M. Saks, and A. Wright. Competitive auctions. *Games and Economic Behavior*, 55(2):242–269, 2006.
- [12] A. V. Goldberg, J. D. Hartline, and A. Wright. Competitive auctions and digital goods. Technical Report STAR-TR-99.09.01, STAR Laboratory, InterTrust Tech. Corp., Santa Clara, CA, 1999.
- [13] Bach Q. Ha and Jason D. Hartline. Mechanism design via consensus estimates, cross checking, and profit extraction. In *Proceedings of the 23rd Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 887–895, 2012.
- [14] J. D. Hartline. Approximation in economic design. Book in preparation, 2012.
- [15] J. D. Hartline and A. Karlin. Profit maximization in mechanism design. In N. Nisan, T. Roughgarden, É. Tardos, and V. V. Vazirani, editors, *Algorithmic Game Theory*, chapter 13, pages 331–362. Cambridge University Press, 2007.
- [16] J. D. Hartline and T. Roughgarden. Optimal mechanism design and money burning. In *Proceedings of the 40th Annual ACM Symposium on Theory of Computing (STOC)*, pages 75–84, 2008. Full version last revised August, 2012.
- [17] J. D. Hartline and T. Roughgarden. Simple versus optimal mechanisms. In *Proceedings of the 10th ACM Conference on Electronic Commerce (EC)*, pages 225–234, 2009.
- [18] J. D. Hartline and Q. Yan. Envy, truth, and optimality. In *Proceedings of the 12th ACM Conference on Electronic Commerce (EC)*, pages 243–252, 2011.
- [19] T. Ichiba and K. Iwama. Averaging techniques for competitive auctions. In *Proceedings of the Workshop on Analytic Algorithmics and Combinatorics (ANALCO)*, pages 74–81, 2010.
- [20] G. Lopomo, L. Rigotti, and C. Shannon. Uncertainty in mechanism design. Submitted, 2009.
- [21] R. Motwani and P. Raghavan. *Randomized Algorithms*. Cambridge University Press, 1995.
- [22] R. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73, 1981.
- [23] N. Nisan. Introduction to mechanism design (for computer scientists). In N. Nisan, T. Roughgarden, É. Tardos, and V. Vazirani, editors, *Algorithmic Game Theory*, chapter 9, pages 209–241. Cambridge University Press, 2007.

A Stochastically Dominating Distributions

Consider a digital goods environment in which the valuation distribution of bidder i is regular and stochastically dominates that of bidder $i + 1$. The optimal auction need not use a monotone price vector, but there is always a near-optimal auction that does. The following result was communicated to us by Dhangwatnotai and Hartline (personal communication, November 2011), and we provide a proof for completeness.

Proposition A.1. *In a digital goods auction with n bidders, if the valuation distribution F_i for bidder i stochastically dominates F_{i+1} for every $i = 1, 2, \dots, n - 1$, and if every distribution F_i is regular, then there is a monotone price vector with expected revenue at least 50% of that of an optimal price vector.*

Proof. We use the probabilistic method. Choose $z \in [0, 1]$ uniformly at random and consider the price vector $\mathbf{p}(z) = (F_1^{-1}(z), \dots, F_n^{-1}(z))$. Since each F_i stochastically dominates F_{i+1} , $\mathbf{p}(z)$ is monotone with probability 1. The expected revenue extracted from bidder i by this random price vector is the expected revenue of a random reserve price p_i drawn from the valuation distribution F_i . Since F_i is regular, the Bulow-Klemperer theorem [5] implies that the expected revenue extracted from the i th bidder is at least 50% times that of a monopoly price; see also [9, Lemma 3.6]. By linearity of expectation, the expected revenue (over z and \mathbf{v}) of $\mathbf{p}(z)$ is at least half that of the optimal auction. There exists a choice of $z \in [0, 1]$ such that the (monotone) price vector $\mathbf{p}(z)$ obtains expected revenue at least half that of an optimal one. \square